IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

PERSONALIZED MEDICINE USING REINFORCEMENT LEARNING

Adarsh Rajput
Course:- M.Tech.

Dr. Harish Dutt Sharma (Assoc. Prof.)
Research Coordinator
Department of Computer Science
Maya Devi University

Abstract: Personalized medicine represents a paradigm shift in healthcare, aiming to tailor medical treatments to the unique genetic, physiological, and environmental characteristics of each patient. Traditional treatment protocols often rely on population averages, which may not effectively address individual variability in disease progression and treatment response. Reinforcement Learning (RL), a branch of artificial intelligence that focuses on sequential decision-making through interaction with an environment, has emerged as a promising approach to optimize individualized treatment strategies. RL algorithms learn optimal policies by continuously evaluating the outcomes of previous decisions and adjusting strategies to maximize long-term rewards.

In healthcare, this reward can represent improved patient outcomes, reduced adverse effects, or cost-effective treatment plans. Recent research demonstrates RL's capability to design adaptive treatment regimens for complex diseases such as cancer, diabetes, sepsis, and psychiatric disorders. These systems integrate data from electronic health records, wearable devices, and genomic databases to generate personalized recommendations. Despite its potential, challenges remain in areas such as data quality, model interpretability, safety, ethical considerations, and clinical adoption. This paper provides a comprehensive analysis of reinforcement learning in personalized medicine, discussing its theoretical foundations, methodologies, real-world applications, advantages, limitations, and future research directions. The study highlights how RL can transform healthcare from reactive treatment models to proactive and adaptive care systems.

Keywords: Personalized Medicine, Reinforcement Learning, Artificial Intelligence in Healthcare, Clinical Decision Support Systems and Precision Treatment Optimization.

1. INTRODUCTION AND BACKGROUND

In the world of medicine, personalized medicine, also known as precision medicine, represents a paradigm shift that focuses on providing tailored medical treatment based on a patient's unique characteristics. In contrast to conventional medicine, which employs a one-size-fits-all approach, personalized medicine utilizes a patient's unique characteristics, such as genetics, molecules, environment, and lifestyle, to identify the key factors that inform an accurate diagnosis, prognosis, and treatment selection. In oncology, cardiology, endocrinology, and neurology, areas where heterogeneity in disease expression and response to treatment is particularly large, this approach has proven extremely useful Advances in biomedical research and the accessibility of high-throughput data contributed to the application of precision medicine in clinical practice. However, the patient data is massive and complicated to analyze; it requires the use of complex computational tools to extract useful insights and make actionable recommendations[1].

The analysis of healthcare data using state-of-the-art technology, such as artificial intelligence (AI) and machine learning (ML), is the payoff. AI systems are currently being used routinely for disease detection, radiological imaging, genomics, and treatment planning, offering unprecedented diagnostic accuracy and efficiency The fact that the AI-driven methods can, in many cases, identify patterns in large and complicated datasets makes them a very suitable tool for applications where we want to integrate and interpret multidimensional data in personalized medicine.

Reinforcement learning (RL) is becoming dominant due to its capacity to learn optimal strategies in dynamic and uncertain environments. RL models improved decision-making over time by taking feedback from the environment and learning interactively. The fact that RL is well-suited to clinical situations where treatment decisions need to adjust to patient responses and disease progression marks a strong advantage in applying RL to optimize health outcomes.

RL is a type of ML motivated by behavioral psychology. It is an agent that interacts with the environment to maximize the cumulative reward by taking actions. The agent, environment, state, action, reward, and a policy mapping states to actions are key components. The agent is a decision-making algorithm, and the environment can include patient characteristics and disease states in medical applications. For example, the reward function typically corresponds to a clinical outcome, such as symptom relief, survival, or normalization of a biomarker. Q-learning is an example of a value-based method, policy gradients are an example of a policy-based method, and actor-critic models are an example of a hybrid method[2]. Deep RL (DRL), which combines deep neural networks with the principles of RL, has been influential in handling complex and high-dimensional healthcare data[3]. These can capture patient dynamics and produce not only personalized treatment strategies but also adapt as their understanding of the patient evolves.

In the context of healthcare, RL can model clinical decision-making as a dynamic process where treatment decisions influence future patient states. By analyzing patient data and outcomes, RL systems can continuously refine treatment policies and adapt to changing patient conditions. This makes RL particularly suitable for chronic disease management and dynamic treatment regimes.

Research indicates that RL can improve personalized treatment planning by learning optimal policies from historical clinical data and simulated patient environments. RL-based clinical decision support systems can outperform traditional rule-based systems by achieving better predictive accuracy and improved patient outcomes.

This paper explores how reinforcement learning can revolutionize personalized medicine by providing adaptive and patient-specific healthcare solutions.

The objective is to gain a thorough understanding of the principles of RL and explore its potential application in personalized clinical treatments within clinical environments. Specifically, this involves exploring the use of RL in multiple medical fields, such as oncology, chronic disease management, and psychiatry, and detailing relevant case studies and real-world examples in these areas where it is applied. Furthermore, upon integrating RL into personalized medicine, the issues and areas of interest that arise from this integration will also be discussed, along with future research directions.

2. OVERVIEW OF PERSONALIZED MEDICINE

2.1. Definition and Evolution of Personalized Medicine

Personalized medicine refers to a modern medical approach that focuses on tailoring healthcare decisions, medical practices, and therapeutic interventions to the individual characteristics of each patient. Unlike traditional medical treatments that follow standardized protocols for large populations, personalized medicine considers variations in genetics, lifestyle, environment, and individual biological factors when designing treatment plans. The primary objective of this approach is to ensure that the right treatment is delivered to the right patient at the right time, thereby improving treatment effectiveness and minimizing adverse reactions.

Personalized medicine, also known as precision medicine, is a new and transformative way of delivering healthcare, where decisions regarding medical interventions and healthcare practices are made on a personby- person basis. While conventional "one-size-fits-all" medicine does not account for differences in people’s genetic makeup, environment, and behavior, personalized medicine does. However, significant advances have been made over the past decade due to developments in genomics, bioinformatics, and systems biology[4]. In oncology, there has been a significant surge in this impact, with the genetic profiling of tumors allowing oncologists to choose therapies based on specific molecular targets. Pharmacogenomics used in neurology is helpful to foretell which group of patients will respond best to antidepressants or antiepileptics. In general, integrating patient-specific biological data enables clinicians to create precise, patient-centric therapies that offer maximum benefit and minimal harm[5].

The concept of personalized medicine has evolved significantly with advances in genomics, biotechnology, and digital health technologies. It integrates patient-specific data such as genetic profiles, molecular biomarkers, clinical history, and environmental influences to predict disease susceptibility and treatment response. Through the use of advanced computational techniques and data analytics, healthcare professionals can identify patterns and develop targeted therapies that are more effective than generalized treatment methods. Personalized medicine is particularly significant in the treatment of complex diseases such as cancer, diabetes, cardiovascular disorders, and neurological conditions, where patients may respond differently to the same therapeutic intervention.

In recent years, the integration of artificial intelligence and machine learning has further enhanced the potential of personalized medicine. Techniques such as reinforcement learning allow healthcare systems to continuously learn from patient data and optimize treatment strategies over time. By analyzing large volumes of clinical and genomic information, these systems can support physicians in making informed decisions and developing adaptive treatment plans tailored to individual patients. Consequently, personalized medicine represents a transformative shift in healthcare from a reactive, disease-centered model to a proactive, patient-centered approach that emphasizes prevention, precision, and improved health outcomes.

2.2. Importance of Personalized Medicine

Personalized medicine has become an essential concept in modern healthcare as it aims to provide medical treatment tailored to the individual characteristics of each patient. Traditional medical practices often rely on standardized treatment protocols that are based on the average response of large groups of patients. However, individuals differ significantly in terms of their genetic makeup, environmental exposure, lifestyle habits, and physiological conditions. These variations can influence how a patient develops a disease and how they respond to specific treatments. Personalized medicine addresses these differences by using patient-specific information to design more effective and targeted healthcare strategies[6].

One of the most significant benefits of personalized medicine is its ability to improve the accuracy and effectiveness of treatment. By analyzing genetic information, biomarkers, and clinical data, physicians can determine which treatments are most likely to work for a particular patient. This approach is especially important in the management of complex diseases such as cancer, cardiovascular disorders, and diabetes, where patients often respond differently to the same therapy. In oncology, for instance, genetic testing of tumors allows doctors to select targeted drugs that specifically attack cancer cells without causing excessive damage to healthy tissues. This leads to improved treatment outcomes and a higher probability of recovery[7].Another important aspect of personalized medicine is its role in reducing adverse drug reactions. Many medications may produce side effects or ineffective results when prescribed without considering individual genetic differences. Pharmacogenomics, a key component of personalized medicine, studies how genetic variations influence a patient's response to drugs. Through pharmacogenomic testing, healthcare providers can determine the appropriate dosage and type of medication that will be safest and most effective for each patient. This not only enhances patient safety but also reduces healthcare costs associated with complications and ineffective treatments[8].

Personalized medicine also contributes significantly to disease prevention and early detection. By identifying genetic predispositions and risk factors, healthcare providers can detect diseases at an early stage or even prevent them before they develop. For example, individuals with a high genetic risk of certain cancers or cardiovascular diseases can undergo regular screening and adopt preventive lifestyle changes. This proactive approach shifts healthcare from a reactive model focused on treating illness to a preventive model focused on maintaining long-term health and well-being[9].

In recent years, technological advancements such as artificial intelligence and machine learning have further strengthened the importance of personalized medicine. These technologies enable the analysis of large volumes of patient data, including electronic health records, genomic data, and real-time health monitoring from wearable devices. Reinforcement learning, a branch of artificial intelligence, plays a crucial role in optimizing treatment decisions by continuously learning from patient responses and adapting medical strategies accordingly. Such systems can assist healthcare professionals in making more accurate and personalized decisions, thereby improving the quality of patient care[10].

2.3. Key Components and Goals

The three areas in which the extension of personalized medicine is genomic medicine, predictive modeling, and tailored therapeutics. With the mapping of the human genome and advancements in next-generation sequencing technologies, genomic medicine has made it possible to identify genetic variations that influence susceptibility to and responses to various diseases and treatments. This is particularly important in oncology, cancer treatment, hereditary diseases, and predicting rates of drug metabolism[11], where genomic information is used to select targeted treatments. Predictive modeling is a form of data analytics or ML that creates models to predict the onset, progression, and potential outcomes of treatments or future disease occurrences. Due to these reasons, these models are used extensively in early diagnosis, risk stratification, patient monitoring, and prevention strategies[12]. The field of personalized therapeutics is much broader than genetic information and also encompasses clinical biomarkers, physiological data, and real-world patient behavior. By doing so, this approach guarantees that interventions are tailored to the individual patient's needs in terms of intensity and timing[13].

2.4. Data Sources in Personalized Medicine

For personalized medicine purposes, data plays an indispensable role, allowing for the use of numerous data sources, ranging from biological and clinical data in the case of algorithms to clinical treatment decisions, which facilitate this. EHRs are longitudinal records of patient data, including demographics, clinical encounters, medications, lab results, imaging reports, and other health IT applications such as care plans, telemedicine, telehealth, clinical decision support, and e-prescribing. However, these datasets are useful for monitoring patient progress over time and gauging patient reaction to treatment. These processes utilize high-throughput platforms that can generate large amounts of data related to genomics, proteomics, metabolomics, and transcriptomics, aiding in the discovery of disease biomarkers and therapeutic targets. The roles of these "omics" profiles are indispensable for personalized treatment[14]. Furthermore, wearables and remote monitoring devices, such as smart watches and fitness trackers, can offer continuous monitoring, providing real-time data on physiological factors (e.g., heart rate, glucose levels) and physical activity metrics. The data streams enable the collection of patient adherence data, recovery trend data, and early signs of deterioration data, providing a further mechanism for offering individualized care[15].

3. REINFORCEMENT LEARNING: AN OVERVIEW

3.1 Definition

Reinforcement Learning (RL) is a branch of artificial intelligence and machine learning that focuses on how intelligent systems learn optimal decision-making strategies through interaction with an environment. In reinforcement learning, an agent learns by performing actions within a particular environment and receiving feedback in the form of rewards or penalties. The primary objective of the agent is to maximize cumulative rewards over time by selecting the most effective sequence of actions. Unlike supervised learning, where models learn from labeled datasets, reinforcement learning relies on trial-and-error experiences to gradually improve decision-making policies[16].

The reinforcement learning framework generally consists of several key components: the **agent**, the **environment**, the **state**, the **action**, and the **reward**. The agent represents the decision-making entity, while the environment refers to the system with which the agent interacts. A state describes the current condition of the environment, and actions are the choices available to the agent. After each action is performed, the agent receives a reward signal that evaluates the effectiveness of the decision. By continuously interacting with the environment and updating its policy, the RL model learns which actions produce the best long-term outcomes[17].

Reinforcement learning has become increasingly important in fields such as robotics, finance, autonomous systems, and healthcare. In the context of healthcare and personalized medicine, RL can be used to design adaptive treatment strategies that evolve according to patient responses and clinical data. By analyzing sequential patient information, reinforcement learning systems can recommend optimized treatment decisions that aim to improve patient outcomes while minimizing risks and side effects[18].

3.2 Components of Reinforcement Learning

Reinforcement Learning (RL) is a powerful approach within the field of artificial intelligence that focuses on learning optimal actions through interaction with an environment. Unlike traditional machine learning models that rely on fixed datasets, reinforcement learning operates through continuous feedback and dynamic decision-making. The learning process in RL is structured around several fundamental components that enable an agent to learn from experience and improve its performance over time. These components include the **agent, environment, state, action, reward, policy, and value function**. Understanding these elements is essential for applying reinforcement learning in complex domains such as healthcare and personalized medicine[19].

i. Agent

The agent is the central decision-making entity in a reinforcement learning system. It is responsible for interacting with the environment, observing the current state, and selecting appropriate actions to achieve a specific goal. In the context of personalized medicine, the agent may represent an intelligent computational model or algorithm that recommends treatment strategies based on patient data. The agent continuously learns from the outcomes of its actions and improves its decision-making policy over time. Through repeated interactions, the agent identifies which actions lead to better long-term results and which ones should be avoided[20].

ii. Environment

The refers to the external system with which the agent interacts. It represents the conditions or circumstances that influence the outcome of the agent's actions. In healthcare applications, the environment can be viewed as the patient's clinical condition, including physiological parameters, disease progression, and response to treatments. When the agent performs an action, such as recommending a specific drug or treatment plan, the environment responds by transitioning to a new state and providing feedback in the form of a reward or penalty. The interaction between the agent and the environment forms the basis of the learning process in reinforcement learning systems[21].

iii. State

A state represents the current situation or condition of the environment at a particular point in time. It contains the information that the agent uses to make decisions. In personalized medicine, the state may include various patient-specific factors such as age, genetic information, medical history, laboratory test results, vital signs, and current medications. By analyzing these variables, the reinforcement learning model can determine the most suitable action for improving patient outcomes. Accurate representation of the state is crucial because it directly affects the quality of the decisions made by the agent[22].

iv. Action

An action refers to the decision or step taken by the agent in response to the current state of the environment. The set of all possible actions is known as the action space. In healthcare applications, actions may include selecting a medication, adjusting drug dosage, recommending a diagnostic test, or modifying a treatment plan. Each action influences the future state of the environment and determines the reward received by the agent. The goal of reinforcement learning is to identify the sequence of actions that maximizes long-term benefits for the system being studied[23].

v. Reward

The reward is a feedback signal that evaluates the effectiveness of an action taken by the agent. It provides a numerical value indicating whether the action produced a desirable or undesirable outcome. Positive rewards encourage the agent to repeat similar actions in the future, while negative rewards discourage harmful or ineffective decisions. In the context of personalized medicine, rewards may represent improvements in patient health, reduced symptoms, decreased hospitalization rates, or lower treatment costs. The reward mechanism is essential because it guides the learning process and helps the agent determine which strategies are most beneficial over time[24].

vi. Policy

A **policy** defines the strategy used by the agent to select actions based on the current state of the environment. In other words, it represents a mapping between states and actions. Policies can be deterministic, where a specific action is chosen for each state, or stochastic, where actions are selected according to a probability distribution. During the learning process, reinforcement learning algorithms aim to discover an optimal policy that maximizes cumulative rewards. In personalized medicine, an optimal policy could correspond to a treatment plan that consistently produces the best possible health outcomes for patients over time[25].

vii. Value Function

The value function is another important component of reinforcement learning that estimates the long-term benefit of a particular state or action. It predicts the expected cumulative reward that the agent will receive by following a specific policy starting from a given state. By calculating value functions, the agent can evaluate different strategies and choose actions that lead to the highest expected reward. In medical decision-making, value functions can help identify treatment options that provide the most favorable long-term health outcomes for patients[26].

viii. Interaction Among Components

The components of reinforcement learning operate together in a continuous cycle. First, the agent observes the current state of the environment. Based on its policy, the agent selects an action. The environment then responds by transitioning to a new state and providing a reward signal. Using this feedback, the agent updates its policy and value function to improve future decisions. Over time, this iterative process allows the reinforcement learning system to learn optimal strategies through experience and adaptation.

In the context of personalized medicine, this framework enables intelligent systems to analyze complex patient data and recommend treatment decisions that evolve according to patient responses. As a result, reinforcement learning provides a powerful tool for designing adaptive and individualized healthcare solutions that can significantly improve patient outcomes.

3.3 Types of Reinforcement Learning Algorithms

Reinforcement Learning (RL) algorithms form the core of intelligent decision-making systems that learn optimal strategies through interaction with an environment. These algorithms enable an agent to determine the best course of action by continuously evaluating the consequences of its decisions. In the field of healthcare and personalized medicine, reinforcement learning algorithms can analyze complex patient data and recommend adaptive treatment strategies that improve health outcomes over time. Various types of reinforcement learning algorithms have been developed to address different learning environments and computational challenges. Among the most widely used algorithms are **Q-Learning, Deep Q Networks, Policy Gradient Methods, Actor–Critic Algorithms, and Proximal Policy Optimization**. Each of these algorithms has unique characteristics and applications in personalized healthcare systems[27].

i. Q-Learning

Q-Learning is one of the most fundamental reinforcements learning algorithms and belongs to the category of model-free learning methods. It enables an agent to learn the value of taking a particular action in a given state without requiring prior knowledge of the environment's dynamics. The algorithm operates by updating a value function known as the Q-value, which represents the expected cumulative reward obtained by performing a specific action in a given state and following an optimal policy thereafter. Over time, the agent learns to select actions that maximize these Q-values, leading to optimal decision-making.

In healthcare applications, Q-Learning can be used to determine the most effective treatment strategies for patients by analyzing historical medical data. For example, the algorithm may learn which drug dosage or therapy sequence produces the best long-term outcomes for a specific disease. Because Q-Learning continuously updates its value estimates based on new experiences, it is particularly suitable for environments where patient responses to treatment evolve over time[28].

ii. Deep Q Networks (DQN)

Deep Q Networks represent an advanced form of Q-Learning that integrates reinforcement learning with deep neural networks. Traditional Q-Learning methods struggle when dealing with large and complex state spaces, such as those found in medical datasets containing thousands of variables. Deep Q Networks address this limitation by using neural networks to approximate the Q-value function instead of storing it in a simple table.

The use of deep learning allows DQN models to process high-dimensional data, including medical images, genomic information, and electronic health records. In personalized medicine, Deep Q Networks can analyze these complex datasets to identify patterns and recommend personalized treatment plans. For instance, DQN models may help optimize chemotherapy schedules by evaluating tumor response and patient tolerance to drugs. This combination of deep learning and reinforcement learning significantly expands the potential applications of RL in modern healthcare systems[29].

iii. Policy Gradient Methods

Policy Gradient Methods represent another important class of reinforcement learning algorithms. Unlike Q-Learning approaches that estimate value functions, policy gradient algorithms directly learn the optimal policy that determines which action should be taken in a given state. These algorithms work by adjusting the parameters of a policy function in a direction that increases the expected cumulative reward.

Policy gradient methods are particularly useful in situations where the action space is continuous rather than discrete. In medical treatment planning, for example, decisions such as drug dosage adjustments or therapy intensity may involve continuous values rather than fixed categories. Policy gradient algorithms can effectively handle such scenarios by learning smooth and flexible policies that adapt to patient-specific conditions. As a result, they are widely used in personalized healthcare systems that require precise and adaptive treatment strategies[30].

iv. Actor–Critic Algorithms

Actor–Critic algorithms combine the advantages of both value-based and policy-based reinforcement learning methods. In this framework, two separate components work together to improve decision-making. The actor is responsible for selecting actions according to a policy, while the critic evaluates those actions by estimating their value. The critic provides feedback to the actor, enabling it to update its policy and make better decisions in the future.

This collaborative structure improves learning efficiency and stability compared to traditional reinforcement learning algorithms. In healthcare applications, actor–critic models can be used to optimize long-term treatment strategies by simultaneously evaluating the effectiveness of medical interventions and adjusting policies accordingly. For example, an actor–critic system may continuously refine treatment recommendations for chronic diseases such as diabetes or hypertension based on patient responses and clinical outcomes[31].

v. Proximal Policy Optimization (PPO)

Proximal Policy Optimization (PPO) is a modern reinforcement learning algorithm designed to achieve stable and efficient policy updates. PPO belongs to the family of policy gradient methods but introduces mechanisms that prevent excessively large policy changes during the learning process. This stability makes PPO one of the most reliable and widely used reinforcement learning algorithms in recent years.

In the context of personalized medicine, PPO algorithms can be used to develop adaptive treatment models that learn from large datasets while maintaining stability and safety in decision-making. For example, PPO-based systems may analyze patient health data to recommend dynamic treatment plans that adjust according to disease progression and patient responses. The ability of PPO to balance exploration and stability makes it particularly suitable for healthcare environments where safety and reliability are critical[32].

vi. Relevance of RL Algorithms in Personalized Medicine

The application of reinforcement learning algorithms in personalized medicine has the potential to transform healthcare by enabling adaptive, data-driven treatment strategies. These algorithms can analyze patient-specific data from electronic health records, genomic databases, and wearable devices to learn optimal treatment policies over time. By continuously updating their decision-making strategies based on patient responses, RL systems can provide highly personalized healthcare recommendations.

Furthermore, reinforcement learning algorithms can assist clinicians in managing complex diseases that require long-term monitoring and dynamic treatment adjustments. Conditions such as cancer, diabetes, and cardiovascular diseases often involve multiple treatment options and uncertain outcomes. RL algorithms can evaluate these options and recommend strategies that maximize patient well-being while minimizing risks and side effects. As research in artificial intelligence and healthcare continues to advance, reinforcement learning is expected to play a crucial role in the development of intelligent clinical decision support systems and precision medicine technologies.

4. WHY REINFORCEMENT LEARNING FOR PERSONALIZED MEDICINE?

RL is applied to personalized medicine because personalized treatment is a sequential decision problem. Caring for patients is a natural sequential decision-making problem in which treatments are changed over time, and each clinical decision influences the future patient states. An example of this is in chronic disease management, oncology, and critical care, where treatment regimens, such as dosage, combination, and timing of medication, often need to be modified to ensure the best response in a patient at a given point in time[33]. However, traditional statistical models are not sensitive to capturing such evolving dynamics over time. In particular, RL is well suited to solve personalized, sequential decision-making problems[34].

In the healthcare environment, there is high uncertainty and variability in how patients respond to treatments, particularly in dynamic and uncertain settings. Patients faced complications in treatment outcomes due to unexpected events and delayed adverse effects. Because of this uncertainty, RL can adapt to changing conditions by continually updating its understanding of the environment. Also, RL algorithms have been applied to determine the optimal timing of starting or stopping ventilator support in a critical care setting and estimate insulin doses based on physiological feedback[35]. In sepsis management, for example, or in corticosteroid therapy[36], RL has been shown to outperform physician-driven policies by identifying adaptive treatments based on observed outcomes. In particular, the flexibility with which RL can handle partial information and delayed feedback is well-suited for clinical scenarios and can be leveraged to develop robust decision support tools that have to adapt as the patient situation evolves.

Patient variability is a significant factor in personalized medicine, and RL also takes it into account. Patients' genetic backgrounds, disease courses, lifestyles, and treatment compliance are different. The RL algorithm can handle this heterogeneity for learning patient- or subgroup-specific policies. Unlike traditional approaches, such as maximizing the average outcome, RL seeks to maximize individual rewards, thus following clinical endpoints like symptom reduction, improved quality of life, or survival[37]. For instance, RL is used to personalize insulin dosing for patients with type 1 and type 2 diabetes, as well as dietary decisions. It has also been shown recently that RL can adjust insulin doses in real-time, depending on meal composition, activity, and glucose levels, which significantly improves glycemic control[38].

Furthermore, RL models have been demonstrated to be able to manage multimorbid conditions by considering the drug interactions and comorbidities[39]. RL has been used in cancer treatment to recommend intermittent androgen deprivation therapy for prostate cancer to achieve the best balance between tumor control and side effects[40]. These diseases are long-term in nature, and treatments are cumulative; therefore, RL's long-term optimization perspective is critical.

Let's now compare RL with other types of learning, such as supervised and unsupervised ones, and see what nice gains we obtain with RL. However, supervised learning has been helpful in diagnostic classification and risk prediction. However, relying on static labeled datasets and being unable to model interventional outcomes has minimized the efficiency of the designed methods. Unsupervised learning is helpful for clustering or identifying hidden patterns in data, but it does not directly impact decisions. Whereas, RL is interventional and adaptive by its essence. It learns from observed data and real-time interactions and recommends and evaluates the consequences of specific clinical actions[41]. This capability makes RL an alternative technology that resides between predictive analytics and action recommendations. Moreover, it offers a dynamic layer for AI-powered personalized medicine by incorporating temporal dependence and decision feedback, which are usually ignored by conventional ML techniques[42].

5. APPLICATIONS OF REINFORCEMENT LEARNING IN PERSONALIZED MEDICINE TREATMENT PROTOCOLS

Figure-1 illustrates the diverse clinical applications of RL in personalized medicine, showcasing its implementation across oncology, chronic diseases, psychiatry, infectious diseases, and rehabilitation settings.

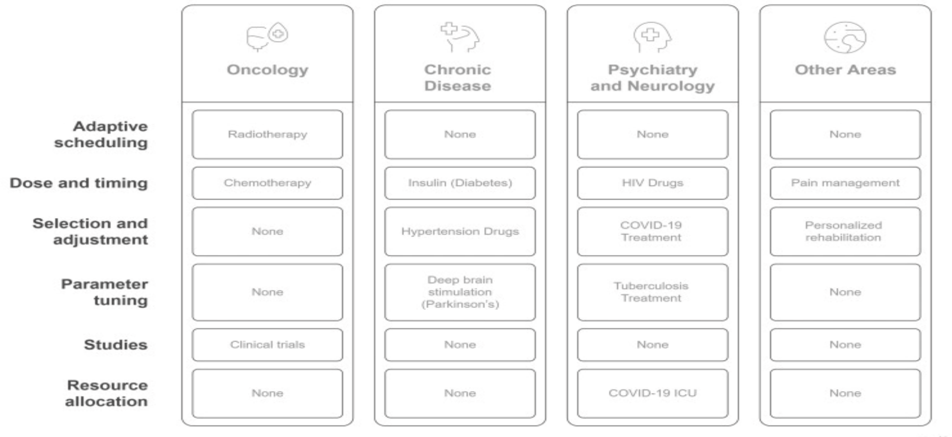

Figure 1. Applications of RL in personalized treatment protocols.

RL: reinforcement learning, COVID-19: coronavirus disease 2019, ICU: intensive care unit.

Oncology: In oncology, treatment personalization is essential due to the range of possible biological tumor types, a patient's genetics, and their response to therapy. Adaptive radiotherapy scheduling is achieved through algorithms in RL that plan the timing and doses of radiation for treatment based on its relationship to tumor growth models, tissue tolerance, and personal response patterns. RL uses optimal policies that adjust based on tumor shrinkage and toxicity data to balance safety and efficacy[43]. Conventional chemotherapy is delivered on a rigid schedule, and there may be over- or under-treatment. RL-based frameworks, such as Q-learning and deep Q-networks, have led to practical case studies and clinical trials for simulating the implementation of chemotherapy protocols and adaptive hormone therapies for prostate and breast cancers. The medical community continues to evaluate these systems for their potential use in actual oncologic decision processes[44].

Chronic Disease: RL is beneficial for learning, as treatment regimens for chronic diseases need to be adapted over time. In particular, real-time insulin dosing for patients with type 1 and type 2 diabetes is a notable application. RL models sample food intake, glucose patterns, and movement-based activities to determine the optimal amount of insulin to inject. Effective glycemic control has been demonstrated in real-world trials using insulin pumps driven by RL or decision support systems[45]. Likewise, in hypertension, where personalized drug selection and continuous dosage adjustment are necessary, RL models have developed optimal treatment policies for lowering blood pressure while minimizing side effects in patients with diabetes and renal impairment[46]. For example, RL models are being trained to personalize how individuals use inhaled corticosteroids based on symptoms, spirometry data, and environmental triggers, aiming to improve patient compliance and reduce medication titration and adherence exacerbations for chronic obstructive pulmonary disease and asthma[47].

Psychiatry and Neurology: These psychiatric and neurological disorders are highly heterogeneous and exhibit a slow response to treatment, making them excellent candidates for RL-based approaches. Traditionally, antidepressant medication personalization utilizes trial-and-error methods for prescribing antidepressants. RL can optimize regimen choices by integrating patient history, genetic markers, and longitudinal mood data to predict the best treatment options for each individual[48]. In epilepsy, RL is applied to personalize the setting of anti-seizure medication or neurostimulation based on predictive models of seizure likelihood, enabling intervention before seizures occur and thus reducing seizure frequency. Additionally, RL extends to deep brain stimulation (DBS) parameter tuning, which involves adjusting stimulation parameters to achieve maximum motor control with minimal side effects, including dyskinesia and cognitive dysfunction, in DBS patients[49].

Infectious Disease: Treatment responses are variable, and resistance patterns of infectious agents evolve, making RL highly applicable. Specifically, in optimizing antiretroviral therapy, RL models develop treatment plans based on patient adherence, viral load dynamics, and the potential for novel mutations leading to drug resistance, aiming to improve long-term viral suppression. RL has been used during pandemics to guide patient-specific ICU interventions, such as selecting ventilation strategies and medications based on dynamic information, including oxygen levels and inflammation markers.

Furthermore, RL strategies have shown benefits in tuberculosis treatment, allowing for the adjustment of treatment regimens to individual pharmacokinetics and adherence patterns, which can improve cure rates in patients with multidrug-resistant tuberculosis. Table-1 illustrates how RL operates in medical settings by comparing different domains based on their decision-making approaches, model types, and achieved outcomes.

Table 1- Applications of RL in clinical domains.

RL: reinforcement learning, DRL: deep reinforcement learning, COPD: chronic obstructive pulmonary disease, COVID-19: coronavirus disease 2019, ICU: intensive care unit.

Domain	Use case example	RL technique applied	Clinical benefit	Reference
Oncology	Adaptive radiotherapy	Deep Q-learning	Dose optimization with minimized toxicity	de Giorgi et al., 2022 [43]
Oncology	Chemotherapy scheduling	Model-based RL	Reduced side effects, improved efficacy	Anzabi et al., 2023 [44]
Oncology	Prostate cancer hormone therapy	DRL	Adaptive hormone suppression	Zheng et al., 2021 [45]
Diabetes management	Insulin pump regulation	DRL	Stable glucose levels	Liu et al., 2024 [46]
Hypertension	Personalized drug titration	Actor-critic models	Controlled blood pressure	Badjatia et al., 2025 [47]
Asthma/COPD	Inhaler adherence optimization	Policy gradient methods	Reduced exacerbations	Munson et al., 2024 [48]
Depression	Antidepressant sequencing	Multi-arm bandit RL	Enhanced treatment response	Sun et al., 2022 [42]
Epilepsy	Seizure intervention timing	Q-learning	Lowered seizure frequency	Kourou et al., 2021 [50]
Parkinson's	DBS parameter optimization	DRL	Improved motor function	Kourou et al., 2021 [50]
COVID-19 ICU	Dynamic treatment planning	Model-free RL	Resource-efficient ICU management	Zheng et al., 2021 [45]

6. DATA SOURCES FOR RL-BASED PERSONALIZED MEDICINE

The effectiveness of reinforcement learning (RL) in personalized medicine largely depends on the availability and quality of healthcare data. RL models require continuous streams of patient information to learn optimal treatment strategies and improve decision-making over time. In personalized medicine, data are collected from multiple sources that capture the biological, clinical, and behavioral characteristics of patients. These data sources enable reinforcement learning systems to analyze complex relationships between treatments and patient outcomes, thereby supporting the development of adaptive and individualized healthcare solutions[50].

One of the most important data sources for RL-based personalized medicine is **Electronic Health Records (EHRs)**. EHR systems contain detailed patient information, including medical history, diagnoses, laboratory test results, prescriptions, imaging reports, and treatment outcomes. By analyzing longitudinal data stored in EHRs, reinforcement learning algorithms can identify patterns in disease progression and treatment responses. These insights allow RL models to recommend treatment strategies that are tailored to the medical history and clinical condition of each patient. EHR data also enable continuous monitoring of patient health, which is essential for developing adaptive treatment policies in personalized medicine[51].

Another significant data source is **genomic and molecular data**, which provide information about a patient's genetic makeup and biological processes. Advances in genomic sequencing technologies have made it possible to identify genetic variations that influence disease susceptibility and drug response. Reinforcement learning systems can integrate genomic information with clinical data to predict how individual patients will respond to specific treatments. This integration is particularly important in fields such as oncology, where genetic mutations play a crucial role in determining the effectiveness of targeted therapies. Genomic data therefore contribute to the development of precision medicine by enabling highly individualized treatment decisions[52].

Medical imaging data also serve as an important resource for RL-based healthcare systems. Imaging techniques such as X-rays, magnetic resonance imaging (MRI), and computed tomography (CT) scans provide visual information about the structure and function of organs and tissues. Reinforcement learning models combined with deep learning techniques can analyze medical images to detect disease patterns, monitor treatment effects, and guide clinical decision-making. For example, imaging data may help RL systems evaluate tumor growth or assess the effectiveness of therapeutic interventions in cancer patients[53].

In addition, **wearable devices and mobile health technologies** are increasingly becoming valuable sources of real-time health data. Devices such as fitness trackers, smartwatches, and continuous glucose monitors collect information about heart rate, physical activity, sleep patterns, blood glucose levels, and other physiological parameters. These data allow reinforcement learning systems to monitor patients continuously and adjust treatment strategies accordingly. Real-time data from wearable devices are particularly useful in managing chronic diseases such as diabetes and cardiovascular disorders, where patient health conditions may change frequently[54].

Furthermore, **clinical trial data** provide high-quality and controlled datasets that are useful for training reinforcement learning models. Clinical trials generate reliable evidence regarding the effectiveness and safety of medical treatments. By analyzing trial data, RL systems can learn from previously tested interventions and apply this knowledge to develop personalized treatment strategies for future patients. These datasets are essential for validating RL algorithms and ensuring that their recommendations align with established medical guidelines and evidence-based practices[55].

The success of RL-based personalized medicine depends on the integration of diverse data sources, including electronic health records, genomic data, medical imaging, wearable device data, and clinical trial information. The combination of these datasets enables reinforcement learning algorithms to analyze complex patient information and generate personalized treatment recommendations. As healthcare data infrastructure continues to improve and advanced data analytics techniques evolve, RL-based personalized medicine is expected to become an increasingly powerful tool for improving patient outcomes and transforming modern healthcare systems.

7. ADVANTAGES OF REINFORCEMENT LEARNING IN PERSONALIZED MEDICINE

Reinforcement Learning (RL) has emerged as a powerful tool in the field of personalized medicine because of its ability to analyze complex data and adapt treatment strategies according to individual patient needs. Unlike traditional machine learning models that rely on static datasets, reinforcement learning operates through continuous interaction with an environment and learns optimal decision-making strategies based on feedback. In healthcare, this environment often represents the patient's health condition, while the actions correspond to medical interventions such as medication, therapy adjustments, or diagnostic procedures. The application of RL in personalized medicine offers several important advantages that contribute to improving patient care and clinical decision-making.

One of the most significant advantages of reinforcement learning is its ability to **develop adaptive treatment strategies**. In personalized medicine, patients may respond differently to the same treatment due to genetic, environmental, or physiological differences. RL models continuously learn from patient responses and adjust treatment policies accordingly. This dynamic learning capability allows healthcare systems to move beyond standardized treatment protocols and develop individualized care plans that evolve over time. As a result, patients receive therapies that are better suited to their unique medical conditions.

Another major advantage of reinforcement learning is its ability to **optimize long-term treatment outcomes**. Many medical decisions have consequences that unfold over an extended period of time. For example, the effects of a specific medication or therapy may not be immediately visible but may influence a patient's health months or years later. RL algorithms are specifically designed to maximize cumulative rewards over time, which means they can evaluate the long-term impact of treatment decisions rather than focusing solely on short-term outcomes. This feature is particularly useful in the management of chronic diseases such as diabetes, cancer, and cardiovascular disorders.

Reinforcement learning also provides an effective framework for **handling complex and high-dimensional healthcare data**. Modern healthcare systems generate enormous volumes of data from sources such as electronic health records, genomic databases, medical imaging, and wearable health devices. Analyzing these diverse datasets can be challenging using traditional analytical methods. RL algorithms, especially when combined with deep learning techniques, can process large-scale data and identify patterns that help clinicians make more accurate and personalized treatment decisions. This ability to integrate multiple data sources enhances the overall effectiveness of personalized medicine.

Another important advantage is the **support it provides for clinical decision-making**. Reinforcement learning can be integrated into clinical decision support systems that assist healthcare professionals in selecting appropriate treatments. By analyzing historical medical data and continuously updating its learning model, an RL-based system can recommend treatment strategies that have shown the highest success rates for similar patients. This support helps physicians make informed decisions while reducing the likelihood of medical errors or ineffective treatments.

In addition, reinforcement learning contributes to **improved efficiency and cost-effectiveness in healthcare systems**. Personalized treatment strategies developed through RL models can reduce unnecessary diagnostic tests, ineffective therapies, and prolonged hospital stays. By identifying the most effective treatment plans earlier in the clinical process, RL systems help healthcare providers allocate resources more efficiently and reduce overall healthcare costs. This economic advantage is particularly important in managing chronic diseases that require long-term care and monitoring.

Finally, reinforcement learning has the potential to **advance preventive and proactive healthcare**. By continuously monitoring patient data and analyzing health patterns, RL models can detect early signs of disease progression and recommend preventive interventions before serious complications occur. This proactive approach shifts healthcare from a reactive model focused on treating illness to a preventive model that emphasizes maintaining long-term health and well-being.

8. FUTURE DIRECTIONS

The integration of reinforcement learning (RL) in personalized medicine is still an evolving field, and future research is expected to further expand its capabilities and applications in healthcare. One important direction is the **integration of large-scale genomic and clinical data** with reinforcement learning models. As genomic sequencing technologies continue to advance, researchers will be able to incorporate detailed genetic information into RL systems to develop highly precise treatment strategies.

This integration will enhance the ability of personalized medicine to predict disease risk, determine drug effectiveness, and design individualized therapeutic interventions.

Another promising area of future research involves the **development of explainable and transparent reinforcement learning models**. Many RL systems currently function as complex "black-box" algorithms, which can make it difficult for healthcare professionals to understand how treatment recommendations are generated. Future research aims to create explainable AI models that provide clear reasoning behind their decisions, thereby increasing the trust and acceptance of RL-based systems among clinicians and patients.

Additionally, the use of **real-time health monitoring technologies**, such as wearable devices and mobile health applications, will significantly enhance RL-based personalized medicine. Continuous patient data collected from these technologies can allow RL algorithms to monitor patient conditions in real time and dynamically adjust treatment plans. This will support proactive healthcare by identifying potential health risks earlier and enabling timely medical interventions.

Another important research direction is the **development of secure and privacy-preserving data-sharing frameworks**. Since personalized medicine relies heavily on sensitive patient data, future research will focus on methods such as federated learning and advanced encryption techniques that allow RL models to learn from distributed healthcare data without compromising patient privacy.

9. CONCLUSION

Personalized medicine represents a transformative shift in modern healthcare by focusing on individualized treatment strategies rather than generalized medical approaches. Traditional healthcare systems have long relied on standardized treatment protocols that are designed for large populations, often overlooking the unique biological, genetic, and environmental differences among individual patients. As medical science advances and the availability of healthcare data increases, the need for more precise and patient-centered treatment methods has become increasingly important. In this context, the integration of reinforcement learning (RL) into personalized medicine offers a powerful framework for improving healthcare outcomes through adaptive and data-driven decision-making.

Reinforcement learning provides a dynamic approach to medical decision-making by enabling intelligent systems to learn optimal treatment strategies through continuous interaction with patient data and clinical environments. Unlike traditional machine learning models that rely on static datasets, RL algorithms are capable of adapting to changing conditions and learning from the outcomes of previous decisions. This ability is particularly valuable in healthcare settings where treatment responses may vary significantly between patients and where long-term outcomes must be considered when selecting therapeutic interventions. By analyzing sequential medical decisions and their effects on patient health, RL models can recommend treatment policies that maximize long-term benefits while minimizing potential risks and complications.

The application of reinforcement learning in personalized medicine has demonstrated promising results in several areas of healthcare, including cancer treatment, diabetes management, intensive care medicine, and mental health therapy. RL-based systems can analyze large volumes of healthcare data derived from electronic health records, genomic databases, medical imaging technologies, and wearable health devices. Through this analysis, these systems can identify complex patterns in patient data and develop individualized treatment strategies tailored to specific patient characteristics. Such capabilities enable healthcare providers to deliver more effective and precise medical interventions, thereby improving patient outcomes and reducing unnecessary treatments or adverse drug reactions.

Despite these significant advantages, the implementation of reinforcement learning in personalized medicine also presents several challenges. Issues related to data quality, model interpretability, patient privacy, and ethical considerations must be carefully addressed before RL-based systems can be widely adopted in clinical practice. Ensuring the transparency and reliability of RL models is particularly important for gaining the trust of healthcare professionals and patients. In addition, the integration of advanced artificial intelligence technologies into healthcare systems requires collaboration among clinicians, data scientists, policymakers, and regulatory authorities to establish appropriate standards and guidelines.

Looking toward the future, continued advancements in artificial intelligence, computational power, and biomedical research are expected to further enhance the role of reinforcement learning in personalized medicine. The integration of genomic data, real-time health monitoring technologies, and explainable AI models will enable the development of more sophisticated and reliable decision-support systems. These innovations have the potential to transform healthcare from a reactive system focused primarily on disease treatment to a proactive system that emphasizes disease prevention, early detection, and personalized care.

Reinforcement learning offers a promising approach to addressing the complexities of modern healthcare by enabling adaptive, patient-centered treatment strategies. By leveraging the power of data analytics and intelligent decision-making, RL-based personalized medicine can significantly improve the quality, efficiency, and effectiveness of healthcare services. As research and technological development continue to progress, reinforcement learning is likely to play a crucial role in shaping the future of precision medicine and advancing the goal of delivering the right treatment to the right patient at the right time.

---

[1] Datta S, Li Y, Ruppert MM, et al.: Reinforcement learning in surgery, https://www.surgjournal.com/article/S0039-6060(20)30825-4/abstract

[2] Loftus TJ, Filiberto AC, Li Y, et al.: Decision analysis and reinforcement learning in surgical decisionmaking. https://www.surgjournal.com/article/S0039-6060(20)30261-0/abstract

[3] Jayaraman P, Desman J, Sabounchi M, Nadkarni GN, Sakhuja A: A primer on reinforcement learning in medicine for clinicians. NPJ Digit Med. 2024, https://www.nature.com/articles/s41746-024-01316-0

[4] Cheng J, Ma X, Li C, et al.: Diet-induced inflammation in the anterior paraventricular thalamus induces compulsive sucrose-seeking. Nat Neurosci. 2022, https://www.nature.com/articles/s41593-022-01129-y

[5] Oh SH, Lee SJ, Park J: Precision medicine for hypertension patients with type 2 diabetes via reinforcement learning. https://www.mdpi.com/2075-4426/12/1/87

[6] National Research Council, *Toward Precision Medicine: Building a Knowledge Network for Biomedical Research and a New Taxonomy of Disease* (Washington, DC: National Academies Press, 2011).

[7] National Cancer Institute, "Targeted Cancer Therapies and Precision Oncology."

[8] U.S. Food and Drug Administration, "Pharmacogenomics and Personalized Medicine."

[9] World Health Organization, Genomics and the Future of Medicine.

[10] National Institutes of Health, "Artificial Intelligence in Precision Medicine Research."

[11] Pelosi AD, Roth N, Yehoshua T, Itah D, Braun Benyamin O, Dahan A: Personalized rehabilitation approach for reaching movement using reinforcement learning. Sci Rep. 2024, https://www.nature.com/articles/s41598-024-64514-6

[12] Mathema VB, Sen P, Lamichhane S, Orešič M, Khoomrung S: Deep learning facilitates multi-data type analysis and predictive biomarker discovery in cancer precision medicine. Comput Struct Biotechnol J. 2023, 21:1372-82. 10.1016/j.csbj.2023.01.043

[13] Eckardt JN, Wendt K, Bornhäuser M, Middeke JM: Reinforcement learning for precision oncology . Cancers (Basel). 2021, 13:4624. 10.3390/cancers13184624

[14] Bologheanu R, Kapral L, Laxar D, et al.: Development of a reinforcement learning algorithm to optimize corticosteroid therapy in critically ill patients with sepsis. J Clin Med. 2023, 12:1513. 10.3390/jcm12041513

[15] Grolleau F, Petit F, Gaudry S, et al.: Personalizing renal replacement therapy initiation in the intensive care unit: a reinforcement learning-based strategy with external validation on the AKIKI randomized controlled trials. J Am Med Inform Assoc. 2024, 31:1074-83. 10.1093/jamia/ocae004

[16] Richard S. Sutton and Andrew G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. (Cambridge, MA: MIT Press, 2018).

[17] Machine Learning literature on reinforcement learning frameworks and decision-making models.

[18] National Institutes of Health, "Artificial Intelligence and Reinforcement Learning Applications in Healthcare Research."

[19] Richard S. Sutton and Andrew G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. (Cambridge, MA: MIT Press, 2018).

[20] Artificial Intelligence research literature on reinforcement learning and intelligent agents.

[21] Machine Learning frameworks describing agent–environment interactions.

[22] National Institutes of Health, research initiatives on AI-driven healthcare systems.

[23] Data Science studies on decision-making models and action spaces in reinforcement learning.

[24] Computational Biology applications of reinforcement learning in biomedical systems.

[25] Computer Science research on policy optimization algorithms.

[26] Mathematical Optimization literature on value functions and decision processes.

[27] Richard S. Sutton and Andrew G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. (MIT Press, 2018).

[28] Q-learning research in reinforcement learning and sequential decision-making models.

[29] Deep Q-Network studies on deep reinforcement learning and neural network-based value function approximation.

[30] Policy Gradient Method literature on policy-based reinforcement learning algorithms.
[31] Actor–critic algorithm research on hybrid reinforcement learning models.
[32] Proximal Policy Optimization modern reinforcement learning frameworks for stable policy optimization.
[33] Yang K, Yu Z, Su X, et al.: PrescDRL: deep reinforcement learning for herbal prescription planning in treatment of chronic diseases. Chin Med. 2024, 19:144. https://link.springer.com/article/10.1186/s13020-024-01005-w
[34] Montemitro C, Ossola P, Ross TJ, et al.: Longitudinal changes in reinforcement learning during smoking cessation: a computational analysis using a probabilistic reward task. Sci Rep. 2024, 14:32171. https://www.nature.com/articles/s41598-024-84091-y
[35] Derbal Y: Can artificial intelligence improve cancer treatments? Health Informatics J. 2022, 28:14604582221102314. https://journals.sagepub.com/doi/10.1177/14604582221102314
[36] Yagishita S: Cellular bases for reward-related dopamine actions . Neurosci Res. 2023, 188:1-9. https://www.sciencedirect.com/science/article/abs/pii/S0168010222002930?via%3Dihub
[37] Engelhardt D, Michor F: A quantitative paradigm for decision-making in precision oncology . Trends Cancer. 2021, 7:293-300.
[38] Mushtaq AH, Shafqat A, Salah HT, Hashmi SK, Muhsen IN: Machine learning applications and challenges in graft-versus-host disease: a scoping review. Curr Opin Oncol. 2023, 35:594-600. https://journals.lww.com/co-oncology/abstract/2023/11000/machine_learning_applications_and_challenges_in.20.aspx
[39] Giordano C, Brennan M, Mohamed B, Rashidi P, Modave F, Tighe P: Accessing artificial intelligence for clinical decision-making. Front Digit Health. 2021, 3:645232. https://www.frontiersin.org/journals/digital-health/articles/10.3389/fdgth.2021.645232/full
[40] Prasad N, Mandyam A, Chivers C, Draugelis M, Hanson CW 3rd, Engelhardt BE, Laudanski K: Guiding efficient, effective, and patient-oriented electrolyte replacement in critical care: an artificial intelligence reinforcement learning approach. J Pers Med. 2022, https://www.mdpi.com/2075-4426/12/5/661
[41] Sun Z, Dong W, Li H, Huang Z: Adversarial reinforcement learning for dynamic treatment regimes . J Biomed Inform. 2023, 137:104244. https://www.sciencedirect.com/science/article/pii/S1532046422002490?via%3Dihub
[42] De Giorgi R, Quinton AM, Waters S, Cowen PJ, Harmer CJ: An experimental medicine study of the effects of simvastatin on emotional processing, reward learning, verbal memory, and inflammation in healthy volunteers. Psychopharmacology (Berl). 2022, 239:2635-45.
[43] Anzabi Zadeh S, Street WN, Thomas BW: Optimizing warfarin dosing using deep reinforcement learning . J Biomed Inform. 2023, 137:104267. https://www.sciencedirect.com/science/article/pii/S1532046422002726?via%3Dihub
[44] Zheng H, Ryzhov IO, Xie W, Zhong J: Personalized multimorbidity management for patients with type 2 diabetes using reinforcement learning of electronic health records. Drugs. 2021, 81:471-82. https://link.springer.com/article/10.1007/s40265-020-01435-4
[45] Liu S, Xu Q, Xu Z, et al.: Reinforcement learning to optimize ventilator settings for patients on invasive mechanical ventilation: retrospective study. J Med Internet Res. 2024, 26:44494. https://www.jmir.org/2024/1/e44494
[46] Badjatia N, Podell J, Felix RB, et al.: Machine learning approaches to prognostication in traumatic brain injury. Curr Neurol Neurosci Rep. 2025, 25:19. https://link.springer.com/article/10.1007/s11910-025-01405-x
[47] Munson BP, Chen M, Bogosian A, et al.: De novo generation of multi-target compounds using deep generative chemistry. Nat Commun. 2024, 15:3636. https://www.nature.com/articles/s41467-024-47120-y
[48] Sun Z, Zeng Y, Yuan T, Chen X, Wang H, Ma X: Comprehensive analysis and reinforcement learning of hypoxic genes based on four machine learning algorithms for estimating the immune landscape, clinical outcomes, and therapeutic implications in patients with lung adenocarcinoma. Front Immunol. 2022, 13:906889. https://www.frontiersin.org/journals/immunology/articles/10.3389/fimmu.2022.906889/full
[49] Kourou K, Exarchos KP, Papaloukas C, Sakaloglou P, Exarchos T, Fotiadis DI: Applied machine learning in cancer research: a systematic review for patient diagnosis, classification and prognosis. Comput Struct Biotechnol J. 2021, 19:5546-55. https://www.csbj.org/article/S2001-0370(21)00428-1/fulltext
[50] Artificial Intelligence research on reinforcement learning and data-driven decision systems in healthcare.
[51] Electronic Health Record systems used for storing and analyzing patient medical data.
[52] Genomics studies on genetic variations and personalized treatment strategies.
[53] Medical Imaging technologies used for disease diagnosis and monitoring.
[54] Wearable medical device health monitoring technologies used for continuous patient data collection.
[55] Clinical trial research studies evaluating medical treatments and interventions.